Interacting with Massive Behavioral Data

Shih-Chieh Su
Qualcomm Inc.
5775 Morehouse Drive
San Diego, CA 92121 USA
shihchie@qualcomm.com

arXiv:1608.07619v1 [cs.LG] 26 Aug 2016

ABSTRACT

In this short paper, we propose the split-diffuse (SD) algorithm that takes the output of an existing word embedding algorithm, and distributes the data points uniformly across the visualization space. The result improves the perceivability and the interactability by the human.

We apply the SD algorithm to analyze the user behavior through access logs within the cyber security domain. The result, named the topic grids, is a set of grids on various topics generated from the logs. On the same set of grids, different behavioral metrics can be shown on different targets over different periods of time, to provide visualization and interaction to the human experts.

Analysis, investigation, and other types of interaction can be performed on the topic grids more efficiently than on the output of existing dimension reduction methods. In addition to the cyber security domain, the topic grids can be further applied to other domains like e-commerce, credit card transaction, customer service to analyze the behavior in a large scale.

CCS Concepts

•Human-centered computing → Visual analytics; Information visualization;

Keywords

data visualization, human interaction, dimension reduction, risk management

1. INTRODUCTION

When there are multiple measures of the each sample, the data is described in the a high dimensional space $\mathcal{H}$ by these measures. To make these high dimensional data points visible to human, a word embedding (or dimension reduction) technique is employed to map the data points to a lower dimensional space $\mathcal{L}$. Usually $\mathcal{L}$ is a two-dimensional (2D) or three-dimensional (3D) space. The word embedding technique of choice attempts to preserve some relationship among the data points in $\mathcal{H}$ after mapping them to $\mathcal{L}$.

For example, the multi-dimensional scaling (MDS) [5] $\mathbb{M}$ tries to preserve the distance between data points, during the mapping from $\mathcal{H}$ to $\mathcal{L}$. The stochastic neighbor embedding (SNE) [6] type of algorithms further emphasize the local relationship ahead of the global relationship. There are other dimension reduction techniques putting emphasis on different favored metrics over relationship. On a specific situation, one particular dimension reduction technique could be more suitable or more efficient than others.

The output from existing dimension reduction algorithm is a set of data points that are non-uniformly scattered around the visualization space, which has some drawbacks:

1. Some data points may overlap with others. Overlap makes the information less perceivable.
2. The data points are denser in some area. The heterogeneity makes human interaction with the data points more difficult.

2. METHODS

In order to better utilize the visualization space, we proposed to distribute the data points evenly over the visualization space. The cloud of data points is deformed in the same space defined by the dimension reduction algorithm of choice. This deformation is denoted as $\mathbb{S}$. In the meanwhile, it is desirable to preserve the point-wise relationship maintained by the dimension reduction algorithm. Our strategy in approaching this goal is prioritized as follows:

1. Points are equally spaced after the mapping $\mathbb{S}$.
2. Point-wise topology is preserved. $\mathbb{S}$ attempts to keep point p_j on the same side of point p_i as before the mapping.
3. Point-wise geometry is loosely followed. When p_i is far from p_j, $\mathbb{S}(p_i)$ is far from $\mathbb{S}(p_j)$.

The algorithm we propose is called the split-diffuse (SD) algorithm (Algorithm 1), which follows the strategies above. In our implementation, the SD algorithm first picks the x-axis as the dimension to split. As in Figure 2 (a), it splits the data points into two groups: the ones smaller than or equal to the median, and the ones larger than the median. Each group goes through this split step again over the y-dimension, as in Figure 2 (b). We recursively split the points in x- and y-dimension iteratively, until there is only one point in current recursion.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

KDD 2016 Workshop on Interactive Data Exploration and Analytics (IDEA'16) August 14th, 2016, San Francisco, CA, USA.

© 2016 Copyright held by the owner/author(s).

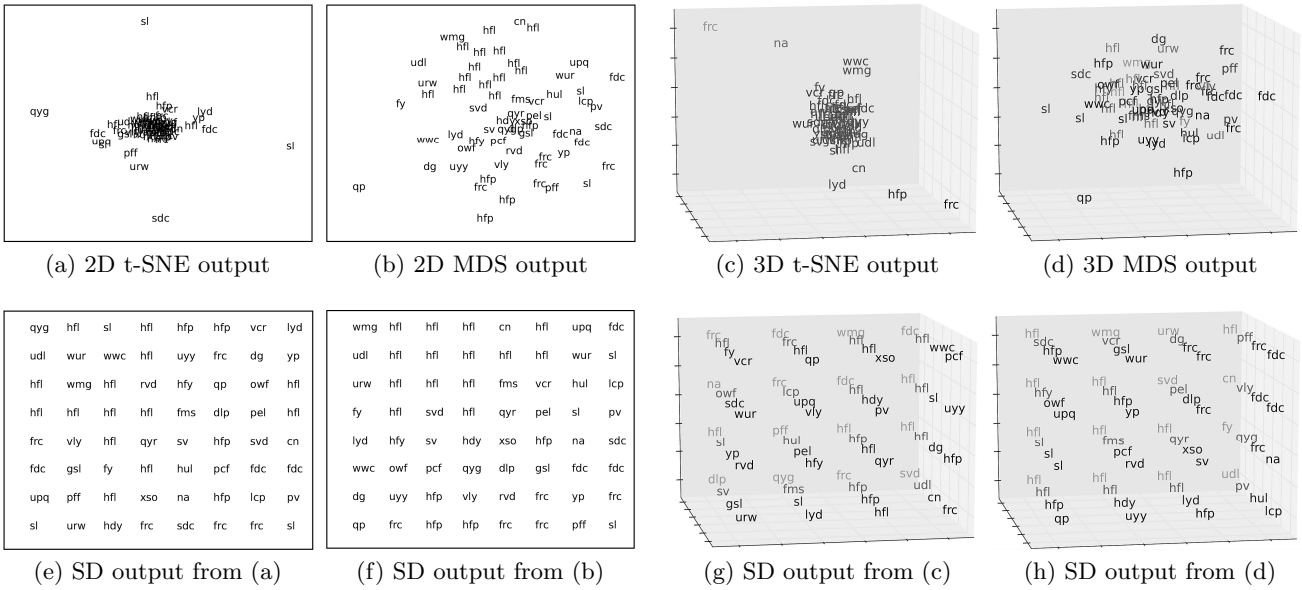

(a) 2D t-SNE output (b) 2D MDS output (c) 3D t-SNE output (d) 3D MDS output

(e) SD output from (a) (f) SD output from (b) (g) SD output from (c) (h) SD output from (d)

Figure 1: The split-diffuse (SD) algorithm takes the output of any dimension reduction technique and distributes the data points evenly while maintaining the topology among them. Example inputs of 64 data points from (a) 2D t-SNE, (b) 2D MDS, (c) 3D t-SNE and (d) 3D MDS are distributed evenly in the same space as shown in (e)-(h) respectively.

Algorithm 1 Split-diffuse algorithm (square of power of 2)

Input: data points $\{p\}$ of length $2^h \times 2^h$, depth $d = 0$, allocation string $c =$ ''

split-diffuse ($\{p\}$, d, c)

  $k \leftarrow$ length of $\{p\}$

  if $k = 1$, **then**

    resolve $\mathbb{S}(p)$ from c

    return p

  end if

  $a \leftarrow mod(\text{depth}, 2)$

  $m \leftarrow$ median of $\{p\}$ in the dimension a

  return ([**split-diffuse** ($\{p : p \leq m|_{dim=a}\}$, d+1, c+'L')],
  [**split-diffuse** ($\{p : p > m|_{dim=a}\}$, d+1, c+'R')])

(a) first-level split (b) second-level splits

Figure 2: The split-diffuse algorithm over the 4×4 layout.

We keep track of the splitting path in string c. At the end of the recursion, the placement each single point p is resolved. The indexes of the SD-mapped points, $\mathbb{S}(p)$, are all integers, and forms a $2^h \times 2^h$ array. This means that the mapped data points are equally spaced in a $2^h \times 2^h$ square. To achieve this uniformity in the space $\mathcal{L}$, the data points are essentially diffused from the denser area to the coarser area by the SD algorithm — hence the name split-diffuse.

Some sample outputs from existing dimension reduction techniques are shown in Figure 1, as well as the corresponding SD outputs. Although we only present the results from t-SNE and MDS, the SD algorithm can be applied to outputs of other techniques such as the principal component analysis (PCA) [3], isomap [4], spectral embedding [1], and totally random trees embedding [2].

3. INTERACTING WITH THE DATA

The motivation to better utilize the visual space comes from the need to interact with massive amount of data. Consider the case that there are millions of items being shipped to a city every month. How can we easily observe the difference in the monthly shipping patterns? When vectoring each item as a data point, putting all the data points on a chart makes the chart hard to read. Instead, using clustering algorithms to group the points and showing the representatives is a better way to present the shipping pattern. Still, with the existing dimension reduction techniques (Figure 1 (a)-(d)), it is difficult to visually compare the difference and interact with the representative points for more detail.

In our use case, we apply the SD algorithm to help analyzing behavioral content in the cyber security domain. The goal of the system is to detect behavioral anomaly based on the access logs. Topics are generated on the content of the logs in a word vector space of $19K+$ dimensions. MDS is applied to reduce the dimension. As shown in Figure 1 (b) and (f), topics are represented by the most relevant keywords, encrypted. Topics close to each other may share the

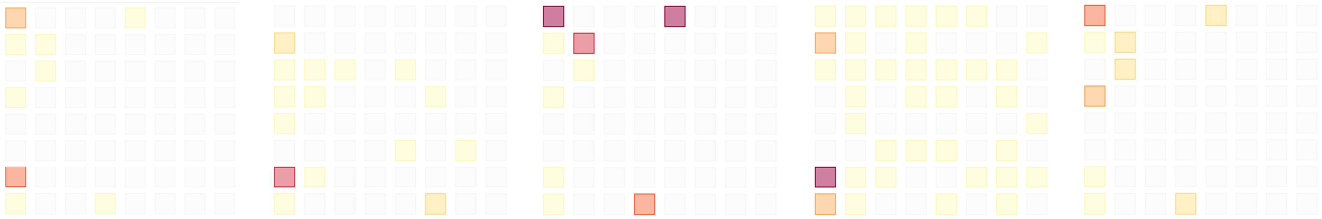

(a) current activities of an entity (b) historical activities of this entity (c) risk against the historical activities of this entity (d) historical activities of the peers (e) risk against the historical activities of the peers

Figure 3: The topic grids. The self risk in (c) is derived from comparing the current activities (a) and the historical activities (b) of a specific entity. The peer risk in (e) is derived from comparing the current activities (a) and the peers' activities (d) of a specific entity.

same representative keyword. The SD algorithm follows to generate the topic grids and visualize different metrics about the behavior of a user (Figure 3).

When not directly displaying the detail keywords about a topic, the topic grids requires less space. At the same time, the human expert still can easily keep track of the topics based on their indexes over all dimensions and compare the difference between different sets of topic grids. Human interaction, which is the ultimate goal of the uniform placement of the data points, can be done more easily on the topic grids than on the raw dimension reduction output as in Figure 1 (a)-(d). For example, the mouse over event on a grid pops up the topical summary, and the click event to overlay the detailed topical activities.

It is also useful to monitor the behavior change over time. In such cases, we reserve a dimension in $\mathcal{L}$ as the time axis. For a 2D space $\mathcal{L}$, a 1D version of SD algorithm is applied to maintain the point-wise topology. The cumulative activities have a shape of curtain. Meanwhile, we can pile up the 2D topic grids on the time axis over the 3D $\mathcal{L}$, as shown in Figure 4. With normal or usual behavior, it is expected to see the consistent hot grids at the same locations over time.

4. FUTURE WORK

In addition to the cyber security domain, the topic grids can be applied to other domains having free-form text logs to analyze the behavior described by the logs. Some possible use cases include e-commerce, credit card transaction, customer service, or others with large volume of behavioral data to be analyzed.

It is also possible to apply the topic grids to the structured data, on which an arbitrary clustering algorithm can generate cluster centers. The data points are then organized into these cluster centers, the same way we use the topic to represent the log entries related to it.

5. REFERENCES

[1] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.

[2] P. Geurts, D. Ernst, and L. Wehenkel. Extremely randomized trees. *Machine Learning*, 63(1):3–42, 2006.

[3] K. Pearson. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.

[4] J. B. Tenenbaum, V. D. Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[5] W. S. Torgerson. Multidimensional scaling: I. theory and method. *Psychometrika*, 17(4):401–419, 1952.

[6] L. Van der Maaten and G. E. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(2579-2605):85, 2008.

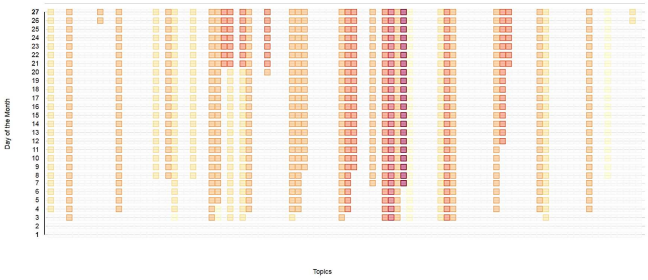

(a) Topic curtain

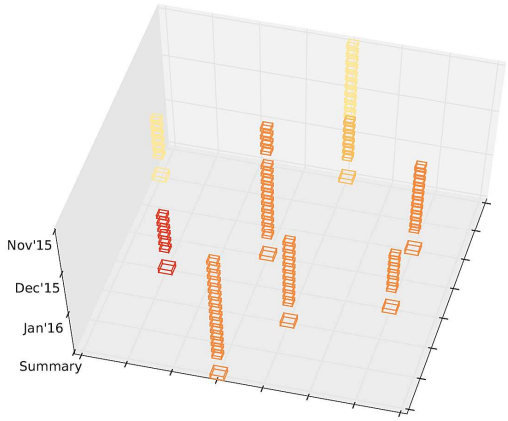

(b) Topic shower

Figure 4: Other formats of the topic grids